

AR-009-454

DSTO-TR-0238

The Tactical Message Interpretation
System

David J. Parker

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

19960429 002

APPROVED FOR PUBLIC RELEASE

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

© Commonwealth of Australia

DTIC QUALITY INSPECTED 1

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

U N C L A S S I F I E D

The Tactical Message Interpretation System

David J. Parker

**Information Technology Division
Electronics and Surveillance Research Laboratory**

DSTO-TR-0283

ABSTRACT

Technical Report

This report describes an experimental computer support system for tactical military intelligence officers. The purpose of the system is to analyse structured military intelligence messages, and to derive conclusions similar to those that might be made by an expert intelligence officer. Details of the structure, functionality and input requirements of the system, and results using information taken from a military training exercise are presented.

APPROVED FOR PUBLIC RELEASE

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

U N C L A S S I F I E D

*Published by DSTO Electronics and Surveillance Research Laboratory
PO Box 1500
Salisbury, South Australia 5108 Australia*

*Telephone: (08) 259 7053
Fax: (08) 259 5619*

*© Commonwealth of Australia 1995
AR-009-454
November 1995*

APPROVED FOR PUBLIC RELEASE

The Tactical Message Interpretation System

EXECUTIVE SUMMARY

This report describes the Tactical Message Interpretation System (TMIS), a system for interpreting formatted intelligence messages about unit sightings. The system was developed in conjunction with the Tactical Military Intelligence Processing System (TMIPS), a technology demonstrator for military intelligence decision support. Techniques developed in TMIPS are expected to be incorporated in AUSTACSS, the Australian Army Automated Command Support System.

The Tactical Message Interpretation System is intended to support the intelligence assessment function of a divisional military headquarters. It reads structured formatted messages about unit and equipment sightings, and estimates the disposition of units on the battlefield, and the enemy's organisational structure, known as the working order of battle (ORBAT). The process extracts information from the intelligence reports, and merges the information into a model of the battlefield, defined in terms of working units and spatial objects. It involves rule-based reasoning using doctrinal, spatial and temporal knowledge, and considers the doctrinal structure, composition and deployment of units, the spatial relationships between sightings, and the sequencing of events. Unidentified units in the model are given arbitrary "working unit" designations, until their actual identities are established. This model is iteratively refined by examining the possible relationships between working units in spatio-temporal neighbourhoods. These relationships are important in identifying or partially identifying units, and in determining the working ORBAT.

The Tactical Message Interpretation System is an object oriented rule based system. It includes a semantic network that represents the relationships between terms used in describing unit sightings. It also includes mechanisms for representing the doctrinal and actual ORBATS, the doctrinal staff tables, and geographic entities. This

information can be loaded for specific scenarios from formatted files. Knowledge about doctrine and tactics are represented by rules, which could be changed to reflect changed circumstances. It was found that an extensive representation of the domain knowledge was required before the system was able to perform any useful function. This suggests that a practical interpretation system is likely to be extremely knowledge intensive, and require a broad and deep knowledge base.

The system has been trialed with messages selected from a mid-level conflict training exercise from the School of Military Intelligence, and generates conclusions similar to those made by a domain expert. It achieved its intended goals, and demonstrates that computer-based systems can be applied to the interpretation and analysis of military intelligence messages, and that decision support tools with interpretative capabilities can be developed for military intelligence analysis. It shows that intelligence analysts could be supported in their tasks, particularly when confronted with heavy message loads and vast geographical areas of interest. However, further work is required to develop a demonstrator with practical utility.

Although the system can interpret some intelligence messages, it was developed as an experimental prototype and should not be considered as a specification for a production system. It operates in a limited domain, and would require further detailed representations of military concepts and knowledge to provide consistent acceptable support. It also anticipates the introduction of messaging systems, such as ADFORMS, for communicating intelligence reports, and the use of structured formatted messages, in contrast to free-text messages. The system has some design problems and unnecessary structural complexity, which make the system less intelligible and more prone to maintenance and update problems. Some of these problems are related to limitations in the expert system shell in which it is implemented, whereas others are related to lack of prior experience in the development environment. It is therefore suggested that any further development of the system be implemented with the latest version of the expert system shell, and that consideration be given to incorporating uncertainty management and truth maintenance.

Authors

David J. Parker Information Technology Division

David Parker graduated Elec Eng (Hons) from the Queensland University in 1968, and was granted a M.Eng and a Ph.D from Newcastle University in 1970 and 1976, respectively. He has worked on analysing and modelling Naval combat systems, and is currently undertaking research into decision support systems for military intelligence analysts.

LIST OF ABBREVIATIONS

AA	Anti-air
AAMG	Anti-air Machine Gun
AD	Air Defence
ADFORMS	Australian Defence Formatted Message System
APC	Armoured Personnel Carrier
Arty	Artillery
AT-xx	a type of Anti-tank Guided Missile
AUSTACSS	Australian Army Automated Command Support System
BDE	Brigade
BN	Battalion
BRDM-xx	a type of Armoured Reconnaissance Vehicle
BTR-xx	a type of Armoured Fighting Vehicle
CLIPS	C Language Integrated Production System
Coy	Company
CSM	Conclusion/Structured Message data structure
DAA	Divisional Assembly Area
DAG	Divisional Artillery Group
DIV	Division
FEBA	Forward Edge of Battle Area
GS	Grid Square
HQ	Headquarters
HVY	Heavy
LT	Light
M-xx	a type of Mortar
MDM	Medium
MBR	Minimum Bounding Rectangle
MR	Motor Rifle (unit)
MSO	Military Spatial Object
mtd	mounted
ORBAT	Order of Battle
PL	Platoon
RAG	Regimental Artillery Group
SAM	Surface to Air Missile
SP	Self-propelled
T-xx	a type of tank
Tk	Tank
TMIPS	Tactical Military Intelligence Processing System
TMIS	Tactical Message Interpretation System

VEH	Vehicle
ZPU-x	a type of towed air defence gun
ZSU-xx-x	a type of self-propelled air defence gun

1. Scope

The aim of this report is to describe an experimental rule-based message interpretation system, and to record some of the technical issues encountered during its development. The purpose of the system is to support tactical military intelligence officers by analysing structured military intelligence messages, and deriving conclusions similar to those that might be made by an expert intelligence officer. This report details the structure, functionality and input requirements of the system, and presents some results using information taken from a military training exercise. Some ideas for further development are raised and discussed.

2. Introduction

The Tactical Message Interpretation System (TMIS) is an experimental computer system for supporting the intelligence assessment function in tactical and operational military headquarters, and in particular, the interpretation of intelligence messages describing events involving enemy units in a mid-level conflict scenario. It is being developed in conjunction with the Tactical Military Intelligence Processing System (TMIPS), a technology demonstrator for military intelligence decision support. Techniques developed in TMIPS are expected to be incorporated in AUSTACSS, the Australian Army Automated Command and Control System. The demonstrator currently supports the entry of messages, conclusions and expectations, the querying and display of message diaries and the intelligence log, the mapping of geographical entities, conclusions and expectations, and the modification and display of the opposing forces organisational structure, called the working order of battle (ORBAT) [Calder 92, Price 92, Milne 92, Lewins 93, Calder 93]. TMIS is a knowledge based tool that interprets intelligence messages and estimates the identity, composition and role of units reported on the battlefield. It also builds and maintains a working ORBAT, in which unidentified units are given arbitrary "working unit" designations, until their actual identity and position in their organisational structure can be established.

The intelligence assessment function is a continuous process involving the interpretation of a stream of intelligence messages, most of which report on the activities of enemy units. The received information is compared and correlated with the available information, and applied to the current estimate of the tactical situation, taking into account the influence that terrain, doctrine, tactics and space/time has on military behaviour. This analysis is currently undertaken by a team of intelligence

officers working cooperatively on a shift duty basis. However, with improved communications and the greater interconnectedness provided by decision support systems, the flow of information will increase to the extent that it will swamp the ability of intelligence staff to perform their tasks effectively without greater levels of support. The purpose of TMIS is to demonstrate that analytical support can be provided for staff tasked to interpret intelligence messages.

TMIS is a knowledge based system in which real-world knowledge is encoded in a format that is both readable by humans and understandable by a computer. real-world knowledge is encoded in a semantic network and a collection of rules. The semantic network encodes knowledge about the concepts, terms and relationships used in describing unit sightings. The rules encode knowledge about entered messages and the disposition and behaviour of units on the battlefield. For instance, the term squad is an instance of type unit size, and is associated to the term platoon by the 'is-smaller-than' relationship.

The primary input to the system consists of intelligence messages. Messages describing events involving units are entered into TMIS in a structured form in which details of the location, time and activity of events, and the characteristics of sighted units, weapons, equipment, vehicles and personnel are entered. The characteristics include the designator, type, size and augmentation of sighted units, and the type, quantity and condition of sighted resources, though entered values may be unknown. This information is associated with one or more units in the working ORBAT, usually unidentified working units, and is then analysed to validate and refine the working ORBAT. In particular unit descriptions and tracks are merged, groupings of equipment to units or elements of units are assigned, superior/subordinate relationships are determined, and units and their roles are identified.

The system is currently designed to support the decision making process, and does not make unilateral decisions about the state of the working ORBAT. If an analysis generates a number of choices, they are preferentially ordered and then presented to the user for a decision. Although the system has not been designed to handle battle indicators, the techniques developed for this system generate information required for some low level indicators, and could be applied in implementing the AUSTACSS "Indicator Expert System" tool.

The prototype has been developed using the CLIPS V5.1 expert system shell [Giarratano 93]. The CLIPS expert system shell is a multiparadigm shell with rule based, object-oriented and procedural programming paradigms. Details of the development environment are given in Appendix B.

3. Knowledge Representation and Input Data

3.1 Structured Messages

Structured messages are formatted representations of free-text military intelligence messages, and are the primary mechanism for entering information into TMIS. The representations are currently limited to information about unit sightings, and in particular to descriptions of events involving sighted military units. Although a sighting is entered into the system as representing a unit and its equipment, the actual event could involve several units (and their equipment), not all of which have been seen. Subsequent processing attempts to resolve this uncertainty by postulating one or more 'actual units' which could explain the sighting.

A structured message has two components, a message header and a message content. The message header component contains the transmission and routing information, such as sender, recipient and time of receipt. This information is not used in the analysis, although the time of receipt is used as a default event time. The message content component contains the unit sighting information, specified in terms of the attributes and relationships of unit, event and location entities. This information is encoded by unit, event and location records, where each record characterises a distinct unit, event or location entity. Units are characterised by their designation, type, size, augmentation and equipment, whilst events are characterised by their location, time, activity and a list of participants. A typical message and its structured representation is illustrated in Appendix A.

3.2 Military Terms

The communication of intelligence messages, as in all forms of communication, requires a common understanding of the universe of discourse by the composers and the interpreters of the messages. In this particular problem, the universe of discourse consists of the terms used for describing events and military units by military personnel. These terms have specific meanings, and in order to interpret the messages, some understanding of the terms must be represented in the message interpretation system. Although the meanings of such terms cannot be fully represented in a computer model, it is possible to encode some real-world knowledge about the concepts and relationships expressed by the terms in a form that is both user readable and understandable by a computer.

3.2.1 Semantic network

In TMIS, knowledge about the concepts expressed by the terms of the universe of discourse is represented by a semantic network in which the nodes represent concepts, and the links represent relationships between them. The network has five distinct

types or classes of concepts, namely unit roles, unit sizes, unit types, unit resources and doctrinal unit types. Each concept is expressed by two or more terms; a descriptive full name, a common abbreviation or brief name, and optionally, a list of aliases. The terms must be unique for each type of concept. For instance, the concept T_MR is a unit type expressed by the full name "Motor Rifle", the brief name "MR", and the alias "Infantry Motor Rifle". Each concept may be involved in one or more relationships, but the involvement of concepts in relationships depends on its type. Details of the relationships for each type of concept in the network are described below. The domain details are taken from the Musorian Armed Forces pamphlet [DoD 80]. It is assumed that any military force can be described in this way.

3.2.2 Unit roles

The role of a unit is its nominal main battlefield function and responsibility, and is one of either combat, combat support or service support. It is not currently used by TMIS in reasoning about intelligence messages, although it is used in structuring ORBAT displays.

3.2.3 Unit sizes

The size of a unit is a measure of its nominal strength, and is one of squad, platoon, company, battalion, regimental artillery group, divisional artillery group, division, corps, army group and unknown. A distinction is made between the sizes of headquarter units and other units. All sizes, except the unknown size, are ranked. All units have a size, although a unit may have an unknown size (an unknown ranking). Because size may have an unknown ranking, the size concept has appropriate definite (e.g. definitely greater than) and indefinite (e.g. possibly greater than) size comparison predicates.

3.2.4 Unit types

Unit types are categories of military units derived from their nominal function or capability, such as anti-air, infantry rifle, motor rifle, surveillance, tank, etc.. This categorisation is very broad, but it is sometimes extended through the use of qualifiers, such as truck mtd, APC mtd, etc., to distinguish between units with the same function but different compositions.

In TMIS, units are categorised into qualified functional types that distinguish between units with different functions and compositions. The qualified functional types form a hierarchy induced by a specialisation relationship over the types. For instance, the unit types "Anti-air (Rifle)" and "Anti-air (Tk)" are specialisations of the unit type "Anti-air". Units of these types have similar anti-air functions, but are specifically structured and equipped for supporting rifle and tank units, particularly

with respect to their mobility and weapon range. An Anti-air (Rifle) company is structured into three Anti-air Machine Gun Platoons (AAMG PL) and one Surface-to-air Missile Platoon (SAM PL), whereas an Anti-air (Tk) Company has two Self-propelled Anti-air Platoons (SP AA PL) and one SAM PL. (The SAM Platoons have different equipment lists, and have different unit types; viz SAM (Rifle) and SAM (Tk).)

The unit type hierarchy allows multiple direct supertypes, although most subtypes only have a single direct supertype. For instance, SAM units are included in both anti-air and air defence units, and have the functionality of both types, and are therefore categorised as a subtype of both anti-air and air defence supertypes.

3.2.5 Unit resource types

Unit resource types are categories of the weapons, vehicles and equipment held by military units, and the personnel staffing military units. They are used to specify the resources held by doctrinal unit types, actual units and working units. Detailed technical characteristics of resource types are not encoded in TMIS.

Unit resource types are similar to unit types in that they can be arranged into a hierarchy induced by a specialisation relationship. The resource type hierarchy also allows multiple direct supertypes and composite types.

Resources can be described as having either atomic or composite structure. For instance, a BRDM/AT3 is a BRDM-2 (an armoured personnel carrier) upon which is mounted an AT/3 (an anti-tank missile). Thus the BRDM/AT3 is described as having a composite structure, and being composed of a BRDM-2 and an AT/3. This structuring is induced by a composition relationship between atomic and composite resource types. Thus the BRDM/AT3 resource type is associated with the BRDM-2 and AT/3 resource types through the composition relationship.

The distribution of resources is encoded within unit resource types. It is derived from and is directly related to the unit holdings (see below). The encoding includes both a nominal and an actual distribution.

The nominal distribution of a resource type is the number of individual items of the resource type, but not of its subtypes or composite types, held by each unit. Abstract generic resource types have no instances, and therefore they do not have a nominal distribution. (It should be noted that an actual individual item may be included in the holdings of many units, because the holdings of subordinate units are included in the holdings of their superior units. The total holdings of a force are given by the holdings of the largest unit.)

The actual distribution of a resource type is the total number of individual items of the resource type, its subtypes and composite types, held by each unit. It is derived from its nominal distribution and the actual distribution of its subtypes and composite types. For example, given that the mortar resource type is an abstract generic type with subtypes 60mm mortar M-57, 82mm mortar M-1937, 120mm mortar M-1943, and 160mm mortar M-160, then the mortar resource has a nominal distribution that is null, and an actual distribution quantifying the number of mortars actually held by each unit, whether they be 60mm, 82mm, 120mm or 160mm mortars. Thus a rifle battalion, which has nine 60mm mortars and six 82mm mortars, is recorded as holding or being issued a total of fifteen mortars.

3.2.6 Doctrinal unit types

Doctrinal unit types are categories of elements of a nominal military force that are identifiable and have a well defined structure and composition. Since doctrinal unit types have specific names which are used to characterise an actual enemy force, they are terms and consequently they are represented in TMIS as a subtype of both the term and unit classes.

Doctrinal unit types are characterised and identified by a unit type and a unit size. (Note that in TMIS, some unit types are subtyped, and have qualified names, to distinguish between similar but distinct unit types. This practise does apply in the field, but to a more limited extent.) In TMIS, the names of a doctrinal unit type are derived by concatenating the names of its type and size.

Each doctrinal unit type has an assigned role, either combat, combat support or service support. A combat unit may have some combat support and service support elements, but not visa versa.

Doctrinal units may be classified into either atomic, composite or functional units. Atomic units may be further subdivided into elemental and HQ units. Atomic units are the smallest identifiable components of a force. They are usually of squad size, but some larger sized units, such as SAM Platoons which are not organised into SAM Squads, are also classified as being atomic. Atomic units are deployed in their entirety, normally with other units to form composite units. Each atomic unit has a defined composition, known as its organic holding, consisting of a specified quantity and type of personnel, weapons, vehicles and equipment.

A HQ unit, such as a rifle corps HQ, is an atomic unit that has operational command over the elements of its associated composite unit. An elemental unit is an atomic unit that is not a HQ.

Composite units are units composed of a HQ unit and one or more subordinate units which are either elemental or composite units. The subordinate units are said to be under the operational command of the HQ unit. For instance, a Motor Rifle Platoon is composed of a Motor Rifle Platoon HQ and three Motor Rifle Squads, whereas a Motor Rifle Company is composed of a Motor Rifle Company HQ and three Motor Rifle Platoons. The organic holdings of a composite unit is identically equal to the the sum of the organic holdings of its component units.

Functional units are groupings of elements of a composite unit that may operate together to perform a task or a function, such as command, combat, reconnaissance, first echelon, etc.. They are not part of the conventional doctrinal ORBAT, and may be structured to give alternative views to the conventional command organisational structure.

3.2.7 Loading the semantic network

The semantic network, which represents knowledge about the concepts expressed by the terms of the universe of discourse, is not predefined in TMIS, and has to be defined and loaded prior to the initiation of the message entry and analysis processes. New terms can be added to the network during the message analysis process, but currently there is no user support for this task.

The semantic network is initially constructed from the term definition file, which is a formatted text file, but the network can be subsequently saved and reloaded as a CLIPS instance file. Each type of term has its own textual format. To simplify the construction of the semantic network, terms used to characterise other terms must be loaded before the terms they characterise. Thus the order in which terms are entered in the term definition file is significant. Firstly, unit and resource types should be loaded in hierachic order from the most generic to the most specific; secondly, unit roles, unit types, unit sizes, and resource types should be loaded before doctrinal units; and thirdly, doctrinal units should be loaded in compositional order from the smallest to the largest.

3.3 Doctrinal ORBAT

The doctrinal order of battle (ORBAT) is a formal documented description of the organisation and composition of a nominal military force. It specifies the operational command structure, unit holdings, staff tables, designation, type and size of each and every unit in the force, and is developed and used as a model for describing an actual military force.

The compositional structure generally corresponds to the command structure, although some functional groupings, such as Divisional and Regimental Artillery

Groups, are included in the hierarchy. Units in the hierarchy are ordered by size, but the ordering is not always sequential. Composite units may be directly composed of units, and in particular specialist units, with various smaller nominal sizes. For instance, a Rifle Division is directly composed of three Rifle Regiments, but also a Reconnaissance Battalion, an Electronic Warfare Company, etc.. As with most hierarchical organisations, the more highly specialised units are small and have shorter lines of command.

The doctrinal ORBAT is not directly represented in TMIS, however the basic structure is represented by the doctrinal unit types. The doctrinal ORBAT can be generated from the doctrinal unit types and unit designation rules.

3.4 Actual ORBAT

The actual ORBAT is a consolidated estimate of the tactical organisation and deployment of an enemy force. The actual structure and composition of the force is guided by, but is not required to follow the corresponding doctrinal ORBAT. Units are modified by the reallocation of equipment, and the attachment and detachment of sub-units from other units or formations of the parent force in response to specific types of operations, likely hostile forces, terrain and operational objectives. The actual ORBAT is derived from information collated from diverse sources, and concentrates on the overall picture and the identity of the deployed formations. It is disseminated down to intelligence units in the field, and is used as a basis for reasoning about the identity of units sighted in their areas of interest.

The actual ORBAT is normally described in terms of identified formations, their compositions and their locations. Each unit is identified by a fully designated name, or if unknown, by an arbitrarily designated name. The name may be qualified or augmented to indicate a significant change to the type or size of the formation. Details of known organisational changes, such as the reallocation of equipment, and the attachment and detachment of units are included in a compositional listing of each unit.

3.4.1 Representing the actual ORBAT

The actual ORBAT is represented in TMIS by the encoded actual unit instances. Each actual unit represents a unit in the actual ORBAT, in particular, its designation, nominal and actual doctrinal unit types, nominal and actual superior, actual subordinates, and its actual and non-organic equipment holdings. Units normally have the same nominal and actual doctrinal unit types, but units with significant non-organic equipment holdings, which may change their operational characteristics, may have actual doctrinal unit types indicating their abnormal operational characteristics.

The command hierarchy is encoded by a superior/subordinate relationship between actual units. Each actual subordinate unit has a nominal superior and an actual superior. The nominal superior is the HQ unit that is specified in the doctrinal ORBAT as having operational command over the subordinate unit, whereas the actual superior is the HQ unit that actually has operational command over the subordinate unit. Normally a unit is under the command of its nominal superior, and the nominal superior is used as the default value for a unit's actual superior. For instance, according to the doctrinal ORBAT [DoD 80], 101 Corps HQ is the superior of 101 Reconnaissance Battalion and 910 Motor Rifle Regiment, but if 910 Motor Rifle Regiment was attached to 301 Rifle Division for an actual operation, then 301 Rifle Division HQ would be the actual superior and be in command of 910 Motor Rifle Regiment.

3.4.2 Initializing the actual ORBAT

The actual ORBAT is initially loaded into TMIS from a formatted text file, but can be saved and reloaded as a file of CLIPS objects. The units are specified in terms of designated doctrinal units, appropriately modified by attached and detached units.

3.5 Working ORBAT

The working ORBAT is an ORBAT that is derived and used solely by staff of an intelligence cell for reasoning about the identities, locations and roles of enemy units sighted or thought to exist in their area of interest. It is similar to an actual ORBAT, except that it records a more detailed, current and localised view of the enemy force. It is derived from the most recently disseminated actual ORBAT and from reports of enemy sightings. Because the reports may contain incomplete, inaccurate and uncertain information, the working ORBAT is more speculative than the actual ORBAT, is subject to re-evaluation and revision, and may have some unidentified units. The unidentified units are arbitrarily distinguished by "working unit" identifiers. The working ORBAT is used solely by staff of an intelligence cell for developing a consistent picture of the tactical situation, and is not disseminated to other intelligence units.

The working ORBAT is initialised by a set of structured messages reporting the locations of those units sighted, identified and included in the actual ORBAT.

3.5.1 Representing the working ORBAT

The working ORBAT is encoded in TMIS by working units and unit descriptions. Working units represent sighted or inferred units in the area of interest. They are described by unit descriptions. Unit descriptions are estimates of the state of an individual unit at a given instance, and are derived from, and are logically supported

by, structured messages and other unit descriptions. Each working unit has a current description and some non-current, replaced descriptions, corresponding to the latest and previous estimated states.

Although each working unit is said to have a current description and some replaced descriptions, there is no direct association between working units and unit descriptions, because the identities of units cannot be determined directly from their sightings. Unit sightings, or more specifically unit descriptions, are associated with abstract units, known as logical units, which are directly associated with working units according to their currently perceived identities. This latter association may change to reflect changes in the perceived unit identities.

The logical units are similar but not equivalent to entered units. Logical units have at least one unit description, but may have a sequence of unit descriptions corresponding to a contiguous sequence of unit sightings. Only those unit descriptions that are known to refer to the same unit are associated with the same logical unit. Unit descriptions that may refer to distinct units are associated with distinct logical units.

3.6 Geographic and Spatial Information

The geographic and spatial information represented in TMIS is currently restricted to named locations and certain spatial and topological relationships, involving unit sightings and military spatial objects. The locations identify and specify the positions of unit sightings and military spatial objects, whilst the relationships are used to express, and in some cases specify, the relative position, orientation and topology between locations, and between unit sightings and military spatial objects.

3.6.1 Military spatial objects

A military spatial object (MSO) is a spatial entity that has tactical significance, such as an assembly area, vital asset, choke point, etc.. They are identified by a name and although they have significantly different characteristics, they are only characterised in the current implementation by their type, location and orientation. MSOs are currently encoded by MSO facts which may be entered either before or during the message entry process.

3.6.2 Locations

Locations are named regions that are used to specify the nominal positions of sighted units and military spatial objects. Each location is identified by a name, and is characterised by a description, a minimum bounding rectangle, and a spatial extent. They may be entered either before or during the message entry process.

Sighted units and military spatial objects do not usually have clearly defined boundaries, and consequently their actual spatial extents cannot be specified accurately. Positions reported in messages are nominal and have varying degrees of accuracy and certainty, ranging from vague geographic names, to more specific four digit grid squares. (Artillery officers usually work to greater accuracy and may report locations using six or eight digit grid references.) Those military spatial objects that can be delineated on a map, generally have complex boundaries which can only be described by a sequence of vectors or map coordinates. Such detail is not usually required for tactical spatial reasoning as long as the spatial relationships between sighted units and military spatial objects are known.

The spatial extent is the set of grid squares (GS) covering the named region. For instance, the locations Burleigh Head and Burleigh, are entered by

(Location (name Burleigh_Head) (GS 4592))

and

(Location (name Burleigh) (GS 4392 4492 4493 4592)),

nominally have spatial extents of GS4592, and GS4392, GS4492, GS4493 and GS4592, respectively. In TMIS, it is assumed that actual spatial extents are included in nominal spatial extents, and that the nominal spatial extents can be used to derive the spatial relationships between sighted units and military spatial objects.

The minimum bounding rectangle (MBR) is the minimum rectangular region aligned to the map grid that includes the named region. It is derived from the spatial extent, if known, and is used to limit the search for unspecified spatial relationships. If the spatial extent is unspecified, then the MBR must be entered as map coordinates. For instance, the location entered as

(Location (name Divisional_Assembly_Area) (xmin 40) (ymin 91) (xmax 45) (ymax 94))

is interpreted as "the location of Divisional_Assembly_Area is bounded by the rectangle aligned to and including grid squares GS4091 and GS4594 (or grid references GR400910 and GR459949)".

3.6.3 Location relationships

A location relationship is a relationship between two locations. They are normally derived from the positional information specified with each location, but they can be entered directly by the user. The currently recognised types of location relationships

are; includes, included in, intersects, north of, south of, east of, west of, north-east of, south-east of, north-west of, and south-west of.

Tactical spatial reasoning is based on relationships between unit sightings, terrain elements and military spatial objects, and in particular their relative positions, and consequently, the actual position and spatial extent of these objects are not necessarily required and do not need to be entered. Reasoning can be based on location relationships entered directly by the user, but because there are relationships between each and every location, there are many more possibly significant location relationships than there are locations. The entry of location relationships should only be considered for locations with positions that cannot be represented by a list of grid squares.

As an example, the fact

(LocationRelationship (location1 Burleigh_Head) (relation INCLUDED_IN) (location2 Burleigh))

can be interpreted as the location of Burleigh Head is included in (the town of) Burleigh. This fact can be directly deduced from the abovementioned location facts.

However, suppose that location Burleigh is included in the location Divisional_Assembly_Area. This information cannot be deduced, with certainty, from the location facts because the actual position of the location Divisional_Assembly_Area is not given. The fact that Burleigh is included in the minimal enclosing rectangular bounded region of the divisional assembly area, only means that Burleigh *may* be included in Divisional_Assembly_Area. The spatial reasoning module will deduce this possibility and seek confirmation from the user. Alternatively, this information may be entered directly by the fact

(LocationRelationship (location1 Burleigh) (relation INCLUDED_IN) (location2 Divisional_Assembly_Area)).

3.6.4 Spatial relationships

Tactical spatial reasoning involves spatial relationships between sighted units and military spatial objects, such as includes, included in, intersects, adjoins, adjacent, at left of, at right of, at front of, and at rear of. Spatial relationships are derived from location relationships, but depend on the orientation of the related objects. They are used by the unit spatial module.

4. Message Entry

The message entry process is a multistaged process for entering structured messages into the message interpretation system, and incorporating the received information into the working ORBAT. It consists of the four interrelated sub-processes; data entry, alignment, association and combination. The data entry and validation processes perform the necessary input and testing of data, read from either keyboard or files, and will not be considered any further here.

4.1 Data entry

The data entry process reads and validates structured messages as directed by the user. The source of the message input stream may be either the keyboard or a data file, and can be redirected after each message or sequence of messages. As mentioned in Section 3.1 and detailed in Appendix I, each structured message consists of a message header and a message body composed of template facts. The message header must precede the message body, but the facts within the message body may be entered in any order.

The entered facts are tested for consistency and referential integrity. Each entered fact consists of a set of named attributes or field values, but only known independent data values need to be entered. The un-entered fields are automatically assigned a derived, unknown or default value according to the extent and values of the entered data. For instance, a location fact has fields for entering the coordinates of its minimum bounding rectangle, but if the extent of the location is entered in terms of grid square (GS) coordinates, then its minimum bounding rectangle does not need to be entered. The minimum bounding rectangle is derivable, and will be derived automatically from the entered data. In contrast, facts with insufficient or inconsistent data or unresolved references are marked invalid and displayed to the user. Messages with insufficient, inconsistent or invalid facts are filtered from the message input stream and discarded. The remaining messages are passed forward to the alignment process.

4.2 Alignment

The alignment process transforms the input source data into a form suitable for correlation and pattern matching. The input data, much of which consists of symbols and strings directly transcribed from free-text military intelligence messages, is converted into a form consisting of standard descriptors or terms. This involves the replacement of each terminological element in the input data by the term in the doctrinal ORBAT with the corresponding name or alias. Terminological elements that are not recognised are rejected by the verification process.

The values of data elements that are used by the correlation and pattern matching processes but are not already explicit in the input data are initialised. Unentered data elements are given default values, and composite data items, such as doctrinal unit type and actual unit identifier, are derived from the component elements.

The output of the alignment process is a set of logical units and logical unit descriptions, where the logical units identify possible distinct sighted units, and the logical unit descriptions describe the involvement of a possible distinct sighted unit in an event.

4.3 Association

The units in the working ORBAT constitute a working model of the enemy force. Support for the model is derived from received information, and in particular, by logical units which represent possibly distinct sighted units.

The initial association of logical units to working units is based solely on unit identifiers. Sighted units are not usually identified, but some may be identified by a fully designated name or a working unit number. Each logical unit has two unit identifiers, the identifier of its sighted unit, such as it may be, and an arbitrary internal identifier for distinguishing it from other logical units. When a logical unit is generated, the identifier of its sighted unit is compared against the identifiers of all the units in the working ORBAT. If the logical unit is found to have a matching unit identifier, then it is associated with the corresponding working unit, otherwise it is associated with a new working unit because it may represent a distinct unit.

4.4 Combination

The logical units associated with a working unit are combined by forming the unit descriptions associated with the logical units into a temporal sequence or track of past locations of what is currently deemed to be the same unit. The temporal sequence is based on the time of observation and not the time of message receipt. The logical unit with the most recent time of observation is known as the current logical unit. Thus a track consists of the unit descriptions of the current logical unit and some merged logical units, all of which are associated with the same working unit. A logical unit is merged into its following logical unit by combining the current unit descriptions of both logical units into a new description, and appending the new description to the following logical unit as its current description. The unit descriptions themselves are combined by selecting the individual independent data elements on a known-for-unknown and a new-for-old basis, and then deriving any dependent data elements from the selected data.

5. Refining the Working ORBAT

The working ORBAT constitute a working model of the enemy force. Units in the working ORBAT derive their support from logical units, but this support may change following the receipt of new information indicating multiple representations of actual units. The working orbat is refined by reevaluating the support for each working unit in the context of the current model, and removing those working units that do not have adequate support. The reevaluation is currently implemented for stationary units and organic resources.

5.1 Stationary units

According to the Musorian doctrine [DoD 80], units are not normally deployed in a stationary position on the battlefield, unless they are preparing for a deliberate attack or undertaking defensive operations. Units are basically on the move, but may be halted in response to basic human needs. In the advance to contact, tactical marches are normally conducted over an eight hour period, with hourly and four hourly rest periods of ten and twenty minutes, respectively. For marches in excess of twelve hours, breaks for meals and sleeping are made at the commander's discretion.

In preparing for a deliberate attack, formations normally move into an assembly area, and remain there only as long as is needed for any necessary regrouping. For a division, this may be one or two days. In defensive operations, and in particular, in area defence, the units are well dug in and are expected to engage the enemy and defend their locations. Although defence is considered a temporary expedient, primarily to gain time to prepare for the counter-offensive whilst preserving their own forces, it could necessitate a lengthy deployment to achieve that aim.

When deployed, units are generally located in close proximity of each other, with first echelon units in front of second echelon units, with respect to the line of advance or the forward edge of the battle area (FEBA). The distance between adjacent units depends on the nature of the terrain and the unit's size, but as a general rule for area defence the gap between combat units is similar to the unit's frontage. For instance, the interunit gaps for infantry rifle platoons, companies and battalions, are up to 300m, 1km and 2km, respectively, in comparison with unit frontages in the range 250-350m, 0.5-1km and 1-2km, respectively [DoD 80]. It should be noted that the interunit gaps of squads and platoons is less than 1km, and consequently, it may not be possible to distinguish between such units, or sightings of such units, from their grid-square locations.

The variation in gap size tends to result in the formation of spatial groupings of units corresponding to their structural groupings. Clusters of smaller units are formed because the gaps between the smaller units are less than the gaps between the larger

units. These groupings may include three, and sometimes even four units of the same type. For instance, a motor rifle company may be located in a group with a similar company, because a motor rifle battalion normally has three subordinate motor rifle companies, whereas a reconnaissance platoon would not be expected to be located in a group with a similar platoon because there are no units with more than one subordinate reconnaissance platoon.

Because of these constraints in the separation and spatial grouping of stationary units, it may be possible to detect multiple sightings from conflicts in the separation and grouping of unit sightings. For instance, units larger than platoons should have distinct non-intersecting grid-square locations, and consequently, sightings of units larger than platoons with intersecting or adjoining locations, are prime candidates for a multiple sighting. A group of sightings, formed by clustering sightings within a specified distance of one another, should be consistent with the unit structure in the doctrinal and, when applicable, the actual ORBAT. A group with inconsistent unit numbers is also a prime candidate of a multiple sighting. However, there are many factors that should be considered in determining which units should be merged to make the group consistent; factors such as co-existing units, or units reported in the same message, spatial separation, and reported characteristics and activities. The criterion currently used is based on co-existence and spatial separation.

5.2 Organic resources

A sighting of military resources, such as weapons, vehicles, equipment and personnel, may be explicitly reported if the observer cannot recognise the sighting as being that of a known unit. The sighted resources must be from one or more units, but in this section we only consider the simple case in which the resources belong to a previously sighted unit. The alternative, more complex case, with sighted resources from many, possibly unsighted units, is considered in section 8. A group of sighted resources is assumed to belong to a sighted unit if it is included in the organic holdings of the unit and is co-located with the unit. In this case, the sightings are merged to maintain the consistency of the working ORBAT.

A group of sighted resources is encoded by a logical unit with a unit description of an unknown unit with some itemised resources. It is initially represented in the working ORBAT by a working unit supported by this logical unit. This unknown working unit will have its own identity, which will distinguish it from other working units. If the sighted resources belong to a sighted unit, then effectively, there will be two representations of the sighted unit in the working ORBAT, one for the unit itself and one for its resources. This type of multiple representation is detected by simply comparing the location and composition of the group of sighted resources with the location and composition of individual sighted units. The owner of the sighted

resources is given by the smallest unit with an intersecting location, and an actual holding, or if unknown, a doctrinal holding which includes the sighted resources.

Once the owner of a group of sighted resources is identified, the logical units for the sighted unit and the group of sighted resources are merged, thereby removing the support for the unknown working unit.

6. Analysis of Sighted Resources

Resources, such as equipment, weapons, vehicles and personnel, are a physical manifestation of a military unit. Military observers are trained to recognize military units from their physical manifestations, and consequently, groups of sighted resources are normally characterised in terms of doctrinal unit types. However, groups of sighted resources that do not conform to a recognised doctrinal unit type are reported by an itemised list of the more significant resources. For example, a message in a training exercise reported the sighting of the following items; 4 ZPU-4, 14 MDM VEH, 6 LT VEH, 8 ZSU-23-4, 10 T-62, 6 BTR-50 and 9 BRDM-2. This group of items was part of the holdings of a Rifle Division of a Rifle Corps, but because it was not a significant component of the Division, it was not described as such.

The purpose of the sighted resource analysis is to estimate the doctrinal type of the units composing a group of military resources. The analysis assumes that the resources are components of the least number of the most specific units, that the resources are specifically held by a formation located near the resources, that the units are elements of the formation holding the resources, and that the composition of units is described by the actual and doctrinal ORBATS.

The analysis consists of two tasks, firstly the selection of the superior formation or unit, and secondly the decomposition of the sighted resources into individual elements of the selected superior unit.

6.1 Selection of superior formation

The superior formation of the sighted resources is required in order to restrict the search for units composing the sighted resources. The system generates a ranked list of candidate superiors based on the assumption that it is more likely that the resources are specifically held by a formation located near the resources. In other words, it is assumed that the resources are not from a more distant formation or from more disparate components of a larger formation. For instance, in the abovementioned example, the resources are assumed to be held by the closest Rifle Division rather the more distant Rifle Divisions or the Rifle Corps.

The list of candidate superior units is initially obtained by comparing the composition of the resource grouping with those of the actual units. The composition of the units in the actual ORBAT is not known in detail, and has to be estimated from their doctrinal types and information about compositional or structural changes, such as non-organic resources, and attached and detached units. The estimate includes all probable subordinate units, in particular, the doctrinal subordinates that are not specifically mentioned in the actual ORBAT. These unaccounted units are initially assumed to have their nominal composition and structure. The composition of each actual unit is aggregated from the composition of its estimated subordinates and its non-organic holding.

The list is then filtered so that only those units that specifically hold the given resources are retained. In other words, if all the resources are held by a subordinate of a superior unit, then the superior unit is removed from the list of candidates. The list of candidates are then ranked according to the spatial distance between the candidate unit and the resource grouping. Currently, the superior unit to be used in the next stage of the analysis is selected by the user from this candidate list.

If the sighted resources are not included in the holdings of an actual unit, then it indicates that the actual ORBAT for the area of operations is incomplete, and that other units have moved into the area of operations. In this situation, the resources are compared with the organic holdings of the known doctrinal units to determine those units with holdings that include the sighted resources. The resources should be included in at least one unit (the top level unit) of the doctrinal ORBAT. For making the selection, the candidate units are ranked by size.

6.2 Decomposition of sighted resources

The decomposition of sighted resources assumes that the resources are from the least number of atomic (vs. composite) units. This assumption is a conservative constraint on the number and size of units composing the sighting. Only those units necessary for composing the sighting are included, except that sufficient numbers of subordinate units may necessitate the inclusion of their corresponding superior units.

Composite units are included in the decomposition if there are sufficient numbers of subordinate units to necessitate the inclusion of their corresponding superior unit. Although the subordinates of a composite units are distinct autonomous units, they normally operate together and function as a single entity. Some subordinates may be detached, but the vast majority of the subordinates remain under the operational command of the composite's HQ. For instance, a Rifle Platoon is normally structured into three Rifle Squads and a Rifle Platoon HQ. A single Rifle Squad may be detached from a Platoon, but the HQ would remain with and in command of the remaining subordinates. Thus a sighting of resources sufficient for two or more Rifle Squads,

would necessitate the inclusion of one or more Rifle Platoons in the decomposition, because they provide more than two thirds of the combat strength of a Rifle Platoon. On the other hand, a sighting of resources sufficient for a Rifle Platoon and a 60mm Mortar Platoon, both of which are subordinates of a Rifle Company, do not necessarily indicate a Rifle Company, because they only provide a minor component of a Rifle Company's combat strength, and could be a detachment in support of some other units.

The decomposition process is based on the distribution or specificity of the sighted resources. Some types of resources are only held by a specific type of unit. These types of resources are called signature equipment because a sighting of this equipment is sufficient to signify or identify the type of the sighted unit. For instance, 60mm mortars are the signature equipment because they identify 60mm mortar squads. Of course, 60mm mortars are also included in the holdings of 60mm mortar squad superiors. The distributions of most types of resources are less specific, although they can still be helpful in identifying types of units. Resources that are distributed to many types of atomic units are more likely to be sighted, but are not good unit discriminators, whereas resources that are distributed to a few types of atomic units are less likely to be sighted, but are good unit discriminators. The specificity of a resource is a measure of the distribution of the resource, and is defined as the number of atomic unit types that include the resource in their holdings. Signature equipment has a specificity of one, whilst the specificity of commonly held items is much greater than one. In the current data set, the specificity of the abovementioned resources are listed in the following table.

Table 1: Specificity of Resources

Resource Type	Specificity
ZPU-4	2
MDM VEH	23
LT VEH	33
ZSU-23-4	2
T-62	4
BTR-50	13
BRDM-2	10

A list of sighted resources is decomposed iteratively by selecting, at each iteration, the most likely unit, and removing its nominal holding from the list of resources. The most likely unit is selected from the elements of the superior formation with holdings that include the most specific resource in the list of sighted resources. In the above example, the first most specific resource would be a ZPU-4 air defence gun, and the units in the candidate set are MR (Truck) Regiment, Rifle Regiment, Arty (Rifle) RAG, AD Company, AA (Rifle) Company, AAMG (MR) Platoon and AAMG (Rifle) Platoon. This candidate set is then filtered to remove those units that do not specifically hold the required minimum number of resource items, namely 4 ZPU-4s, with the result that, in this example, the initial candidate set is reduced to an AAMG (Rifle) Platoon. If necessary, the units in the candidate set are ranked by correlating the nominal resource holdings of each unit with the current list of sighted resources. The current list of sighted resources is then amended by subtracting, with a lower bound of zero, the resource holdings of the selected unit from the resource list.

The decomposition process is then reiterated with the amended list of resources, viz 10 MDM VEH, 5 LT VEH, 8 ZSU-23-4, 10 T-62, 6 BTR-50 and 9 BRDM-2 (Note that the number of MDM VEHs and LT VEHs have changed because an AAMG (Rifle) Platoon has four MDM VEHs and one LT VEH.), and continued until there are no resources left in the list. In this example, the next most specific resource is a ZSU-23-4 air defence gun, and the units in the candidate set are Tk (Rifle) Regiment, MR (Truck) Regiment, Arty (MR) RAG, AD Company, AA (Tk) Company, SP AA (Tk) Platoon and SP AA (Rifle) Platoon. The complete list of resources is decomposed into five units; AAMG (Rifle) Platoon, AA (Tk) Company, Tk (Rifle) Company, Recon (Tk) Company, and Tk Regiment Cmd.

7. Structuring the Working Units

The working ORBAT, like the doctrinal and actual ORBATS, has a hierachic organisation induced by the compositional relationship between military units. Sightings may identify units which, although nominally distinct, belong to different levels in the hierarchy and are compositionally related. Any analysis of working units has to take into consideration the likely hierarchy, and the fact that units sighted and reported at one level may compose units sighted and reported at a higher level.

The compositional relationship between working units is not directly observable, but in some circumstances, it can be deduced from the location and behaviour of the units. As a general rule, the members of a unit move together in a coordinated manner, and are deployed in neighbouring locations. All things being equal, the distance between units is in proportion to their size, and consequently, units tend to form spatial clusters that correspond to the grouping of units in the hierachic structure. This clustering of

working units can be used to estimate the compositional relationship between units, and the actual disposition of the enemy force.

In TMIS, the compositional hierarchy of the working ORBAT is determined by two complementary incremental processes, the first estimates the most direct superior of each current working unit, whilst the second aggregates spatial groupings of units into composite working units. The processes are executed whenever the working ORBAT is modified, and use the following assumptions;

- All working units are part of a single compositional hierarchy, which implies that the hierarchy consists of a single topmost level unit and a collection of subordinate units. The top level unit is taken from the current actual ORBAT or, if not specified, the doctrinal ORBAT.
- The working ORBAT is derived from the doctrinal ORBAT with a minimal level of restructuring, particularly with respect to combat elements. Units are normally redeployed as functional entities with their support elements.
- A working unit may have a generic unit type, and consequently may be one of many specific unit types.
- The composition of working units conforms to their doctrinal types unless specifically indicated by their decomposition or size augmentation.
- The working ORBAT may be incomplete, in the sense that the direct superior of some subordinate working units may not be included in the current working ORBAT.

Note that the working ORBAT is only an estimate of the organisation and composition of sighted units and their superiors. Units are not specifically decomposed unless one or more of their subordinates are involved in individual sightings.

7.1 Estimating the most direct superior

The most direct superior of each current working unit is estimated on a unit by unit basis. The process involves the generation of a set of candidate superiors, removing conflicting candidates, and then ranking and selecting the most likely candidate. The candidate superiors are generated by rules of the following form. (Note that these rules are not CLIPS rules, which have a different syntax and incorporate procedural as well as declarative conditions. See Rule 1 in Appendix D. Also note that the indefinite and definite articles are used to indicate unbound and bound variables, respectively.)

If

a *subordinate-working-unit* has-superior an *unknown-superior*, and
the *subordinate-working-unit* has-doctrinal-type a *subordinate-doctrinal-type*,
and
the *subordinate-doctrinal-type* has-kind a *doctrinal-kind*, and
a *working-unit* has-doctrinal-type a *doctrinal-type*, and
the *doctrinal-type* has-element an *element*, and
the *element* has-doctrinal-type the *doctrinal-kind*,

Then

the *working-unit* is a candidate superior of the *subordinate-working-unit*.
(Rule 1)

If

a *subordinate-working-unit* has-superior a *superior*, and
the *subordinate-working-unit* has-doctrinal-type a *subordinate-doctrinal-type*,
and
the *subordinate-doctrinal-type* has-kind a *doctrinal-kind*, and
a *working-unit* has-higher-superior the *superior*, and
the *working-unit* has-doctrinal-type a *doctrinal-type*, and
the *doctrinal-type* has-member a *member*, and
the *member* has-doctrinal-type the *doctrinal-kind*,

Then

the *working-unit* is a candidate direct-superior of the *subordinate-working-unit*.
(Rule 2)

Rule 1 is applicable for working units with an unknown superior (the initial default value), and selects any working unit that, according to the doctrinal ORBAT, has elements of the appropriate kind. Rule 2 is applicable for working units with a known superior, and selects any subordinate of that superior that, according to the doctrinal ORBAT, has members of the appropriate kind. There are similar rules for other situations, such as fully designated working units with compositions specified in the initial actual ORBAT, and composite units of neighbouring HQs.

The selection of candidate superiors is on a unit-by-unit and also a rule-by-rule basis, which may create conflicts in the candidate set. These conflicts are removed by retracting the conflicting candidates by rules of the form;

If

a *working-unit* is a candidate superior of a *subordinate-working-unit*, and
a *superior-working-unit* is a candidate superior of the *subordinate-working-unit*, and
the *superior-working-unit* is not a candidate direct-superior of the *subordinate-working-unit*, and
the *superior-working-unit* has-element the *working-unit*,

Then

retract the *superior-working-unit* is a candidate superior of the *subordinate-working-unit*.

(Rule 3)

If

a *working-unit* is a candidate direct-superior of a *subordinate-working-unit*,
and
the *working-unit* is a candidate superior of the *subordinate-working-unit*,

Then

retract the *working-unit* is a candidate superior of the *subordinate-working-unit*.

(Rule 4)

Rule 3 removes high level indirect superiors when more direct superiors within the line of command are candidates, whereas Rule 4 ensures that direct superior relationships are known as such.

The candidate superiors of each subordinate unit are then ranked according to the distance between their locations, and placed in an ordered list from which the most likely candidate is selected. The selected compositional relationships are then incorporated in the working ORBAT by modifying the appropriate unit descriptions.

7.2 Aggregating spatial groupings

The second process used in constructing the composition hierarchy of the working ORBAT aggregates spatial groupings of individual units into composite working units. The process is initiated by groupings of working units that indicate the location of unreported composite units. The process involves the generation of a set of candidate groupings which, if confirmed by the user, are used to identify and locate new composite working units. Once identified, the composite units are incorporated into the working ORBAT. The composite units are then compositionally associated with their components by the process described in the previous section.

The candidate groupings are generated by rules of the form;

If

a *working-unit* has-doctrinal-type a *doctrinal-type*, and
the *working-unit* has-role combat, and
the *working-unit* has-activity deployed, and
the *doctrinal-type* has-most-likely-superior a *superior-doctrinal-type*, and
the *superior-doctrinal-type* has-at-least-two-members-of-type the *doctrinal-type*, and
the *doctrinal-type* has-kind a *doctrinal-kind*, and
an *another-working-unit* has-doctrinal-type the *doctrinal-kind*, and
the *another-working-unit* has-role combat, and
the *another-working-unit* has-activity deployed, and
the *working-unit* is-next-to the *another-working-unit*,

Then

the *working-unit* is-in-candidate-grouping-with the *another-working-unit*.
(Rule 5)

The grouping identified by this rule is restricted to neighbouring deployed combat units of similar doctrinal types. The groupings confirmed by the user are then used to identify and locate new composite working units.

8. Spatial Reasoning

Spatial reasoning during intelligence message analysis occurs both at a low or general level and at a high or domain level. Reasoning about locations and the spatial relationships between locations occurs at a low level, and is not concerned with specialised military knowledge. Reasoning about the spatial relationships between and among units and military spatial objects occurs at a high level, and is primarily driven by military knowledge. In this section we only consider the low level reasoning.

8.1 Low level spatial reasoning

8.1.1 Relationships between locations

A number of relationships between locations can be derived from the details of their position and spatial extent. The topological relations "includes", "included in" and "intersects" are derived from the grid square data, whilst the directional relations "north of", "south of", "east of", "west of", "northeast of", "southeast of", "northwest of", and "southwest of" are derived from the bounding rectangle data. The distance between locations is also derived from the bounding rectangle data.

Relationships between locations that are not characterised by grid squares are initially derived from the bounding rectangle data. The topological relations derived from this method are then confirmed by the user.

The directional relationships are derived from the orientation of the vector between the centre points of the bounding rectangles. Each of the eight directions is defined by a non-overlapping sector of 45° centered on the nominal direction.

8.1.2 Relationships between units and military spatial objects

Some additional directional relationships between units and military spatial objects can be derived from their orientations in conjunction with the directional relations between their locations. The relative directional relations "left of", "right of", "front of" and "rear of" are used in domain level spatial reasoning. In TMIS, each of these relative directions is defined by a broad overlapping sector of 135° centered on the nominal direction.

8.2 Domain level spatial reasoning

Domain level spatial reasoning is not highly developed in TMIS. It is currently restricted to rules involving the formation of assembly areas, and the actual roles of units deployed about and within assembly areas and the FEBA (forward edge of battle area), and the membership of echelons. It includes rules of the form;

If

a *working-unit* has-activity deployed, and
 the *working-unit* has-location a *unit-location*, and
 a *military-spatial-object* has-type assembly_area, and
 the *military-spatial-object* has-location an *assembly-location*, and
 not the *working-unit* is-spatially-included-in the *military-spatial-object*, and
 the *unit-location* is-near the *assembly-location*,

Then

the *working-unit* is-spatially-included-in the *military-spatial-object*.

(Rule 6)

If

a *working-unit* has-activity deployed, and
 the *working-unit* is-spatially-included-in a *military-spatial-object*, and
 the *military-spatial-object* has-type assembly_area, and
 the *working-unit* is-at-rear-of the *military-spatial-object*,

Then

the *working-unit* has-current-role defence.

(Rule 7)

If

a *working-unit* has-activity moving, and
a *military-spatial-object* has-type FEBA, and
the *working-unit* is-at-front-of the *military-spatial-object*,

Then

the *working-unit* has-current-role reconnaissance.

(Rule 8)

If

a *working-unit* has-activity deployed, and
the *working-unit* has-superior a *superior-unit*, and
a *military-spatial-object* has-type assembly_area, and
the *working-unit* is-spatially-included-in the *military-spatial-object*, and
the *military-spatial-object* is-spatially-included-in the *superior-unit*, and
the *working-unit* is-at-front-of the *military-spatial-object*,

Then

the *working-unit* has-current-echelon first.

(Rule 9)

Rule 6 defines a mechanism for extending the spatial extent of an assembly area to cover a build-up of units. Rule 7 and 8 determine the current roles of units which may differ from their nominal doctrinal roles. Rule 9 concerns the positioning of the echelons of the dominant unit of an assembly area. Although these rules are simplistic and may need to be revised, they do give a flavour of the type of information that can be inferred from the locations of units and military spatial objects. The use of this information in identifying sighted units has yet to be considered.

9. Further Developments

9.1 Military spatial objects

The current TMIS implementation makes no representational distinction between types of military spatial objects (MSOs) such as assembly areas, vital assets, observation points, avenues of approach, and choke points. They are currently represented by a data structure (a non-ordered or deftemplate fact) that specifies their name, type, location and orientation. Although the spatial relationships between these objects may be directly represented or determined from their locations and orientations, other relationships, such as the "can observe" relationship between observation points and avenues of approach, and specific features, such as the significance of a vital asset, the effective width of an avenue of approach, etc., are not represented. By representing military spatial objects by more specific data structures, more specific relationships and features could be represented.

In representing military spatial objects, a distinction needs to be made between persistent terrain-based objects and transient unit-based objects. The former could be represented in a geographical information system, and then downloaded by and copied in the message interpretation system. Unit-based objects, such as "assembly areas" and "forward edge of battle areas" are more appropriately represented by groupings of units, rather than spatial regions, although they could still be depicted by regions on battle maps.

9.2 Spatial templates

Spatial templates are annotated graphical representations of unit deployments. They are included in the doctrine of a force, and are used as a guide for the positioning and the movement of units during training exercises and tactical operations. A template is a scaled mapping of a unit deployment with graphical symbols representing deployed units. The symbols are from a defined unit symbology that represents unit type and unit size.

Although the primary purpose of spatial templates is to act as a guide for the positioning and movement of units during training exercises and tactical operations, they can also be used by an observer in recognizing units and in predicting the positioning and movement of units. In particular, knowledge of the spatial templates of a hostile force may be used in analysing sightings of elements of that force.

Intelligence staff use spatial templates during message analysis to gain further information about sighted units and the state of operational preparedness. They visually compare groups of mapped unit sightings with the applicable spatial templates to determine the best match, and hence the most likely deployment. They then use the most likely deployment to predict further sightings (if the deployment is incomplete), and to reclassify current sightings (if their unit type is unknown or generic).

The application of spatial templates to the recognition of units has not been considered in this implementation. The matching of unit sightings to spatial templates by machine is a resource intensive process involving the selection of candidate unit sighting configurations, the positioning (rotation and translation) of templates, and possibly the rubber-banding (distortion) of templates to allow for the impact of terrain.

9.3 Unit groupings

Besides the command hierarchy organisation, units are also organised into functional groupings and associations, such as forward detachment, advance guard, first echelon, mounted units, supported units, etc.. The groupings are formed as and when required, to perform specific roles or attain specific objectives. The breakdown and composition

of these groupings varies according to the situation, but some guidelines and constraints are included in military doctrine. The inclusion of functional groupings and associations in the doctrinal knowledge base may improve unit recognition.

9.4 Temporal templates

Templates describing the build-up and timelines of military operations are not represented in the current system. These are used by intelligence staff to determine the state of operational preparedness, and to estimate the launch time of operations. The current system does not distinguish between temporal and inferred states, and although it does support some temporal reasoning, the representation of unit states would need to be modified to enable this form of temporal reasoning.

9.5 Truth maintenance

The inferencing process in TMIS is frequently confronted with the necessity to select among plausible alternatives. For instance, there may be alternative associations between entities and the input data, or alternative conclusions drawn from incomplete or inaccurate data. These sets of alternatives can be handled in different ways. The system currently handles each set of alternatives by either selecting the most likely alternative, or requesting the user to make a selection. The selected alternative is then added to the current context, and the inferencing process continues.

Although this algorithm is simple to implement, it has some problems. The algorithm maintains a single state and only allows one solution to be considered at a time. This is reasonable when there is always one alternative that is significantly more plausible than the others, but it is unsatisfactory when the significant alternatives are equally plausible. It is extremely difficult to compare solutions for two or more alternatives, and may not necessarily lead to a solution that is consistent and satisfactorily explains all the input data.

An alternative approach is to incorporate a justification-based truth maintenance system [Doyle 1979]. These systems have a dependency-directed backtracking algorithm that searches for a consistent solution by retracting inferences, backtracking through the data dependencies, and selecting a different alternative. However this algorithm could consume considerable resources finding a solution consistent with the input data, particularly since the input data is being revised continuously by the inclusion of new messages.

Another approach is to incorporate an assumption-based truth maintenance system [De Kleer 1986a]. In this instance, an assumption corresponds to a selected alternative. These systems essentially maintain multiple contexts where each context consists of all the data derivable from a consistent set of assumptions. Each set of alternatives can

be represented by a oneof disjunction, thereby greatly reducing the number of consistent sets of assumptions, but each oneof disjunction of size n will still increase the number of consistent contexts by a factor of n (vs 2^{**n} for n independent assumptions). Whilst there are efficient techniques for labelling the assumption sets of derived data, this technique become intractable with large numbers of independent assumptions. However, the number of contexts could be pruned by considering their plausibility.

9.6 Fuzzy reasoning

Fuzzy reasoning is a form of reasoning that provides a framework for solving imprecisely formulated problems by handling sources of uncertainty which are not well defined. It is based upon the theories of fuzzy sets and fuzzy logic. Fuzzy sets are used to represent the uncertainty due to the lack of clarity or completeness in terminology. The fuzzy logic provides a simple mechanism for combining evidence, and enables the inferencing of possibly imprecise conclusions from sets of possibly imprecise premises.

Most of the information used in message analysis is imprecise and incomplete. The identity, position, type and size of sighted units are imprecise and incomplete. The doctrinal statements about the structure, composition and operational behaviour of forces are generalisations that, by its very nature, have only a vague resemblance with reality. Reasoning with this type of information usually involves the consideration of a number of competing hypotheses. By explicitly representing uncertainty, the fuzzy logic may be used in combining information and in establishing the level of support for any inferred conclusions or hypotheses.

A fuzzy version of CLIPS has been developed by the National Research Council of Canada which handles fuzziness and uncertainty [Riley 95]. This version may provide better support for fuzzy reasoning.

9.7 Object oriented development

Whilst CLIPS 5.1 supports the concepts of objects, classes and message passing (message handlers), the object component of the system is not fully integrated with the rule based inferencing engine. Instances of user-defined classes cannot be pattern-matched on the left-hand side of rules, so patterns involving objects cannot be used as antecedents or conditions of rules. Any information that is to be used as antecedents of rules, such as messages, sightings and conclusions, must be stored in the factlist and cannot take advantage of the object oriented features. Composite or structured information is decomposed into unstructured flat facts, which increases not only the complexity of the application, but also the maintenance costs. Although the factlist can be viewed, it is generally very lengthy and is only viewable as a complete listing

which, besides changing whenever a rule fires, requires significant processing resources to generate and list in a readable form.

A new version of CLIPS has been released (CLIPS 6.0 with updates 6.1 and 6.2) which has a significantly improved rule/object integration and the construction of modular programs. It allows instances of user-defined classes to be pattern-matched on the left-hand side of rules, and as a result the information represented by objects, including responses from their associated message handlers, becomes directly accessible in the inferencing process. The new module construct partitions the knowledge base and can be used to implement functional blackboards. These features could greatly improve the TMIS design, and simplify both system development and maintenance.

The object oriented extensions could also be used to improve the inferencing control mechanism in TMIS. Control of the inferencing process is a major problem in TMIS. Forward chaining production systems are controlled by inbuilt conflict resolution control strategies. These control strategies are appropriate for some problem solving tasks, but are inappropriate for other tasks. They assure search exhaustivity, but do not explicitly restrict the problem search space to contexts that satisfy the problem constraints. This may not be an issue for problems with small search spaces, but problems with large search spaces become intractable unless the search space is restricted to contexts satisfying the problem constraints. It is also an issue for systems that interact with users, because each interaction should be focussed on a particular context. Whilst these contextual restrictions can be represented in the fact base, it usually results in obscure and convoluted rules that are a major source of maintenance problems. For instance, TMIS iterates through the input message stream, with each iteration focussing on specific problem solving contexts. The refocusing on contexts can create a cycle in the search space causing unforeseen and undesired side effects. This is not a problem in some reasoning systems which segregate the problem solving aspects of the system from its control aspects. Truth maintenance features, such as de Kleers' consumers and schedulers [de Kleers 1986], have the ability to enforce exhaustivity, satisfy constraints and maintain contextual focus. By implementing these features in the CLIPS object oriented language, it may be possible to eliminate the control problem in the rule base.

10. Conclusions

This paper describes an expert system for interpreting messages in support of the tactical and operational intelligence assessment function. The system interprets structured intelligence messages about sighted units and equipment. Messages entered into the system are analysed to determine the disposition and identity of units in an area of interest, and to generate and maintain a working ORBAT of the enemy

force. The analysis takes into consideration the location and timing of events, the composition and structure of the doctrinal and actual ORBATs, and the doctrinal guidelines for the deployment and placement of units.

Although most intelligence messages handled by the Australian Defence Forces currently contain unstructured free text, there are significant resources being applied, both in Australia and overseas, to the development of structured messaging systems. The purpose of these systems is to support the construction, validation, transmission and distribution of structured messages by users in the field, and to reduce or eliminate the need for unstructured free text messages. The incorporation of such systems into the defence forces, would greatly facilitate the application of message interpretation systems which could be configured to receive messages directly from the communication network.

The practical utility of the TMIS message interpretation system is constrained by the depth and breadth of its knowledge concerning the representation and content of information conveyed by intelligence messages. The knowledge represented in the system is focussed on the categorisation, composition and deployment of military units, and has a narrow scope and limited practical utility. Although the knowledge represented in TMIS was derived from the hypothetical Musorian Armed Force [DoD 80], it is not restricted to this force. In TMIS the details encapsulated in the doctrinal ORBAT are dynamically loaded and not explicitly encoded, so the knowledge base can be tailored to represent forces with similar characteristics by loading it with the appropriate doctrinal ORBAT. However, this tailoring is constrained to forces with similar characteristics, in particular, conventional regular armed forces. These forces may differ in their details, but they have similar types of entities and relationships. On the other hand, irregular insurgency forces are less constrained and have varying entities and relationships, reflecting their greater autonomy and obscure, adaptable organisational structures. Systems interpreting messages involving these forces need to be adaptable and require models that are specifically updated and maintained by users in the field. TMIS does not support this level of adaption, although the development environment itself does support change. The current version of TMIS can interpret information about the ORBAT of a regular armed force, and would require additional knowledge for interpreting other types of information.

The current emphasis by Australian Defence Forces on the management of short warning time conflicts and the provision of rapid deployment forces, raises the possibility of Australian forces being deployed with little prior knowledge of the actual enemy force. Formal message interpretation system are not likely to be effective under these circumstances, until knowledge of the actual enemy force is acquired and represented, and a knowledge base is constructed.

Whilst the TMIS demonstrator has shown that AI techniques can be effectively applied to the interpretation of military intelligence messages, it also raised some issues that need to be addressed in future demonstrators. The knowledge base, as it is currently structured, is difficult to develop and maintain in a consistent and reliable manner because it has compiled within it much of the control over the inferencing process. This also means that the knowledge base is obscure and difficult for domain experts to understand and incorporate new knowledge. Later versions of the CLIPS expert system shell have provided features that may be useful in handling this issue. The issues of uncertainty and consistency are not addressed in the TMIS demonstrator. Since the information upon which the interpretations are based are typically inaccurate, incomplete, sparse, irrelevant, untimely, conflicting, etc., the inferences generated from it should be qualified by some degree of uncertainty and inconsistency. Techniques have been developed for handling these issues in other, more constrained domains, but whether they can be successfully applied to this domain is an open question. These issues will need to be considered in any future prototype development.

The CLIPS expert system tool which was used to development the TMIS demonstrator has both advantages and disadvantages. It provides a complete environment for the construction of rule and/or object based expert systems, and is both a productive development and delivery expert system tool. It has a large international user base within the military, government, corporate and academic communities, and has an active internet news group. It has an ongoing development program with the phased release of upgrades and new versions. The source code is written in the C language and is freely available. It can be linked to external functions or embedded in other applications. It is also used as the root of other third-party products, such as FuzzyCLIPS and DynaCLIPS. The tool does not directly support uncertain reasoning and truth maintenance, though it does support logical dependencies. It has a very primitive user interface, however third-party graphic user interfaces that can be linked to the tool are available. Given that this tool is still being enhanced both by the original developers and third-party groups, it should continue to be an excellent tool for future prototyping work.

11. References

Calder, P., (1992) "Modelling Conclusions in Intelligence Assessment", ITD - 92 - 23, November 1992.

Calder, P., (1993) "Working Orbats in Intelligence Assessment", ITD - 93 - 15, July 1993.

Department of Defence (Army Office) (1980) "Manual of Land Warfare part three volume 2 pamphlet No. 2 The Musorian Armed Forces" January 1980.

Doyle, J., (1979) "A Truth Maintenance System", Artificial Intelligence, Vol. 12, p 231-272, 1979.

De Kleer, J., (1986a) "An Assumption-based TMS", Artificial Intelligence, Vol. 28, p 127-163, 1986.

De Kleer, J., (1986b) "Problem Solving with the ATMS", Artificial Intelligence, Vol 28, p 197-224, 1986.

Giarratano, (1989) "Expert Systems - Principles and Programming", PWS-Kent, Massachusetts, 1989.

Lewins, N., (1993) "A Heuristic Parser for Military Sighting Reports", ITD - 93 - 23, September 1993.

Milne, P., Milton, S., Smith, J. L., (1992) "Geographic Object-Oriented Databases - A Case Study", Journal of Geographic Information Systems, 1992.

Monai, Francesco, F. M., Chehire, T., (1992) "Possibilistic Assumption based Truth Maintenance System, Validation in a data Fusion Application", Thomson-CSF/RCC, France, 1992.

Orchard, B., (1994) "FuzzyCLIPS 6.02", KSL, National Research Council of Canada, 1994.

Price, R., Calder, P., (1992) "The Expectation Model and Template Matching for Intelligence Processing", ITD - 92 - 22, November 1992.

Gary Riley, (1995) "The CLIPS FAQ", Internet - CLIPS ListServer, April 1995.

Appendix A

Fact based representation of Intelligence Messages

Structured messages are formatted representations of free-text military intelligence messages, and are the primary mechanism for entering information into TMIS. The information is expressed in the form of template facts which are then entered into the system. When expressions in the form of template facts are read by the data entry system, they are immediately asserted into the CLIPS fact list as template facts, and become available for processing.

The representational form of structured messages was developed for unit sighting analysis, and is restricted to information about unit and equipment sightings. Each message is assumed to describe a temporal sequence of events involving one or more units or groupings of equipment. This information is expressed in the form of CLIPS template facts, in particular, Entered-CSM, Entered-unit and Entered-event facts. Some messages may also include geographic information that can be expressed in the form of Location and LocationRelationship facts. This information can be entered either before or within a structured message.

A structured message has one Entered-CSM fact followed by one or more Entered-unit and Entered-event facts. For example, the message

"2 x RIFLE BN HAVE CROSSED TERANORA WATERS AND
DEPLOYED IN AREA OF GS4582, GS4583, GS4780 AND GS4781. 4 x
BRDM-2 WITH AT-3 DEPLOYED WITH THE BN."

received by 31 BDE from 3 DIV at 260730 May, is represented by the CLIPS facts

```
(Entered-CSM (type MESSAGE) (time 260730May) (from "3 DIV") (to "31 BDE"))

(Entered-unit (serial 1) (type rifle) (size bn))

(Entered-unit (serial 2) (type rifle) (size bn))

(Entered-unit (serial 3) (type UNKNOWN) (size UNKNOWN) (equipment 4 "BRDM-2 with AT-3" OK))

(Entered-event (serial 1) (activity MOVING) (location Teranora_Waters) (participants 1 2 3))

(Entered-event (serial 2) (activity DEPLOYED) (location GS4582_GS4583) (participants 1))

(Entered-event (serial 3) (activity DEPLOYED) (location GS4780_GS4781) (participants 2))

(Entered-event (serial 4) (activity DEPLOYED) (location GS4582_GS4583_GS4780_GS4781) (participants 3))

(Location (name Teranora_Waters) (xmin 48) (xmax 48) (ymin 80) (ymax 80))
```

The Entered-CSM fact is the header or envelope of the message. It has slots for the originator, recipient and time of receipt, although only the time of receipt is interpreted, and this is only used as a default for unentered event times.. Although structured messages are the primary mechanism for entering information into TMIS, the data entry system is more generic and also allows users to enter conclusions which have a form similar to messages. These different types of information are distinguished in the Entered-CSM fact (Conclusion or Structured Message) by the type slot which can have the value of MESSAGE or CONCLUSION. The template for Entered-CSMs is specified by the construct:-

```
(deftemplate Entered-CSM
  (field type
    (type SYMBOL)
    (default CONCLUSION)
    (allowed-symbols MESSAGE CONCLUSION))
  (field time
    (type LEXEME INTEGER)
    (default NOT-ENTERED))
    ;date/time group of message receipt
  (field from
    (type LEXEME)
    (default NOT-ENTERED))
    ;originator's unit
  (field to
    (type LEXEME)
    (default NOT-ENTERED)))
    ;addressee's unit
```

The Entered-unit facts corresponds to distinct unit entities described in the message. Each Entered-unit fact is assigned a distinct serial numbers which is used to identify it as a participant in one or more events. It also has slots for entering attributes of the unit such as its designation, type, size and augmentation. The equipment slot is used to specify a distinct grouping of equipment, which may otherwise be unrecognised. It consists of a list of count/type/status triples. The Entered-unit template is specified by;

```
(deftemplate Entered-unit
  (field serial
    (type INTEGER SYMBOL)
    (default NOT-ENTERED))
    ;arbitrary serial for identifying participants of events
  (field designator
    (type LEXEME)
    (default NOT-ENTERED))
    ;unit designator e.g. "?/1/301"
  (field type
    (type LEXEME)
    (default NOT-ENTERED))
    ;unit type e.g. rifle, RIFLE
  (field size
    (type LEXEME)
    (default NOT-ENTERED))
    ;unit size e.g. div, DIV, "div HQ"
  (field size-augmentation
    (type LEXEME)
    (default NOT-ENTERED)
    (allowed-symbols NOT-ENTERED)
    (allowed-strings "(?)" "(+)" "(-)" "(0)"))
    ;unit augmentation
  (field working-unit
    (type INTEGER SYMBOL)
    (default NOT-ENTERED))
    ;numeric component of working unit number. The symbol
    ;NEW causes the next available WU number to be
  allocated
    (allowed-symbols NEW NOT-ENTERED)
    (default NOT-ENTERED))
  (multi-field equipment)
    ;list of count / type / status triples
```

Working units numbers are arbitrary designators used by the intelligence staff to refer to unidentified units known to be located within their area of interest. They normally have the format WUnn, but only the numeric component is used here. The equipment field should only be entered if it is known that the equipment belongs to the equipment holdings of the specified unit, otherwise any grouping of equipment should be associated with an unknown unit.

The Entered-event facts are associated with distinct event entities. Each Entered-event fact is assigned a distinct serial number which establishes the temporal sequence of events when the event times are not given. Events are characterised by their location, time, activity and a list of participants. Details of the template are as follows.

```
(deftemplate Entered-event
  (field serial
    (type INTEGER SYMBOL)
    (allowed-symbols NOT-ENTERED)
    (default NOT-ENTERED))
  (field activity
    (type SYMBOL)
    (allowed-symbols NOT-ENTERED UNKNOWN DEPLOYED MOVING RECON PATROL
AMBUSH SIGHTING)
    (default NOT-ENTERED))
  (field location
    (type SYMBOL)
    (default NOT-ENTERED))
  (field orientation
    (type SYMBOL)
    (allowed-symbols NOT-ENTERED UNKNOWN NORTH SOUTH EAST WEST NORTH_EAST
SOUTH_EAST NORTH_NORTH_WEST SOUTH_WEST)
    (default NOT-ENTERED))
  (field time
    (type LEXEME INTEGER)
    (default NOT-ENTERED))
  (multi-field participants
    (default NOT-ENTERED)))
  ;serial of event within message
  ;reported activity
  ;reported location, either a symbolic location, a GS, a GR
  ;an underscored sequence of GRs or GSs
  ;e.g. TeranoraWaters, GS1234, GS1234_GS1235, GR123456
  ;orientation of deployment or movement
  ;date/time group of event
  ;serials of participating units
```


Appendix B

Development Environment

The TMIS message interpretation system was developed using the CLIPS V5.1 expert system shell which consists of the following components :-

- Functions: These are pieces of executable code, identified by a specific name, which returns a useful value or performs some useful side effects.
- Fact List: The fact list is a database of assertions, commonly called the working memory. Facts are one of the basic high-level forms for representing information in CLIPS. A fact consists of a sequence of fields that can be dynamically added or removed from the fact list. They are stored in one of two formats: ordered or non-ordered. Ordered facts consist of an unconstrained sequence of fields separated by spaces and enclosed in parentheses, where each field may be an instance of any primitive data type. No restrictions are placed on the ordering of fields in ordered facts, however the ordering is significant for the pattern matching process. Non-ordered or deftemplate facts consist of a set of named fields or slots specified by a user-defined template. The template specifies the type, default value and allowable range of values for each field.
- Knowledge Base: The knowledge base is a collection of production rules, commonly called the production memory. Rules are one of the primary methods for representing knowledge in CLIPS. Each rule consists of an antecedent and a consequent. The antecedent is a set of conditions or patterns involving facts in the fact list which must be satisfied for the rule to be placed on the agenda. The consequent is a set of actions that are to be performed when the rule is fired. It normally specifies the information that is to be added to, or removed from, the working memory.

This version of CLIPS (in contrast to version 6.0) does not directly support the structuring of the knowledge base into modules or rule sets, although some form of structuring can be established in terms of rule conditions. For instance, a rule set for the "advance to contact" "phase of battle" could be specified by placing the condition

(phase-of-battle advance-to-contact)

in each rule of the rule set.

- Inference Engine: Performs the basic execution cycle consisting of the pattern matching, conflict resolution, rule execution phases. The pattern matching phase, matches the antecedents of each rule against the current state of the fact list, and determines whether the rule is to be activated or deactivated. Activated rules, namely those rules whose antecedents are currently satisfied, are placed on the agenda as a rule activation, which is a rule together with the facts that match its antecedents. Deactivated rules are removed from the agenda.
- Agenda: The agenda is the list of all the rules which have their antecedents satisfied and have not yet been executed. It behaves like a stack in the sense that

new rules are placed on the stack, and the top rule is the first one to be executed. The placement of a newly activated rule on the agenda is determined by the rule saliency, and the current conflict resolution strategy. The CLIPS inference engine provides seven conflict resolution strategies (depth, breadth, simplicity, complecity, LEX, MEA, random), from which the current strategy is dynamically selectable. None of the conflict resolution strategies are definitive, and there is a certain level of arbitrariness in the actual placement of rules on the agenda. Rules are deactivated (removed from the agenda) either when they are executed or when their antecedents become unsatisfied. In contrast to procedural programs, the inference engine has overall control over program execution (the firing of rules), but procedural control can be established by the use of appropriate procedural facts and conditions.

- **Classes and Objects:** A class is a system or user-defined template for the common properties and behaviour of the objects which are instances of that class. The common properties and behaviour are defined in terms of slots or named fields, and message-handlers or procedural code, respectively. Slots are placeholders for values associated with instances. Each instance of a given class has the same set of slots, but the values bound to each slot is only associated with that instance. In this programming paradigm, an object responds to a message sent to it by executing the corresponding message-handler on its slot values, and returning the appropriate result.

All classes, except the system defined Object class, have one or more superclasses, from which they inherit properties and behaviour. Conflicts between multiple superclasses are resolved by the class precedence list. The slots and message-handlers explicitly specified by a class augment or override those specified by its superclasses.

- **Interpreter:** CLIPS programs are normally interpreted by the CLIPS interpreter, but they can be compiled into run-time modules which can be executed with less memory and processing resources. An interpreted CLIPS program, which consists of the constructs loaded in the CLIPS shell, such as functions, templates, rules and classes, can be dynamically modified by the addition or removal of individual constructs during the interpretation process. Such modifications could be triggered by changes in the input data stream, and might represent learning or the gaining of new knowledge. In contrast, the executable of the run-time modules is non-dynamic, and cannot be modified by the addition or removal of individual constructs.

Appendix C

TMIS Functional Architecture

The TMIS functional architecture consists of the following components :-

- a) **Control Module:** This module provides the user interface and manages the data entry and the message analysis subprocesses. It enforces a phased analysis of the messages, and ensures a stable context and uniform interface for user interaction with the subprocesses.
- b) For developmental purposes, the processing is controlled by the user through a menu interface. Messages can be selectively input from the keyboard or data files, and then selectively processed.
- c) **ORBAT Module:** This module consists of some classes of objects for representing military units and equipment. The objects represent information about the structure, composition and roles of doctrinal unit types and actual units. The objects also represent the semantic relationships between different types of units and items of equipment. They can be used for describing new units, and in identifying sighted units. The module includes the domain specific classes UnitSize, UnitType, UnitRole, Designator, ResourceType, DoctrinalUnitType and ActualUnit.
- d) **Spatial Module:** This module determines the spatial relationships between locations of unit sightings and military spatial objects from their actual or default bounding rectangles. Requests from the user for geographic information, such as the distance by road between locations, are made when and if required.
- e) **MessageEntry Module:** This module links together the information entered by the user in the form of Entered-message, Entered-unit and Entered-event facts, generates sighted unit descriptions, and associates the sighted unit descriptions with units in the working ORBAT.
- f) **LogicalUnits:** This module matches and links elements of the sighted unit descriptions to objects in the doctrinal ORBAT, updates the unit composition facts and maintains the consistency between the symbolic and object representations. This linking enables the matching of doctrinal information in the context of the sighted units.
- g) **Unit Experts:** These are collections of rules or knowledge bases that encapsulate human expertise or knowledge concerning a particular aspect of unit data analysis. TMIS currently has four unit experts, viz: MergeLogicalUnits, AnalyseSuperior, AnalyseEquipment and UnitSpatials, which are described below. The unit experts suggest changes to the unit descriptions, and if accepted, replace the current unit descriptions with new descriptions. The retention of replaced descriptions enables backtracking and the formulation of explanations.

- h) MergeLogicalUnits: This unit expert merges matching sighted units that may be generated by multiple sightings of the same event, resulting in multiple messages.
- i) AnalyseSuperior: Uses knowledge of the doctrinal ORBAT and the locations of known sighted units to suggest possible direct or indirect (element of) superiors for units with unknown superiors. Establishes appropriate links if the solution is accepted.
- j) AnalyseEquipment: Uses knowledge of equipment allocations and locations of known sighted units to suggest the units indicated by the sighted equipment. Replaces the equipment sighting by one or more sighted units.
- k) UnitSpatial: Uses knowledge of doctrine and locations of sighted units and military spatial objects to establish the role and echelon of sighted units.

Appendix D

Production Rule Syntax

Rules are defined in CLIPS using the defrule construct which has the following syntax.

```
(defrule <rule-name> [<comment>]
  [<declaration>] ;Rule Properties
  <conditional-element> ;Left-Hand Side
  =>
  <action> ;Right-Hand Side
```

Single valued and multi-valued variables have names of the form ?aSymbol and \$?anotherSymbol, respectively. Similarly, single valued and multi-valued wild-cards are indicated by a ? and a \$?, respectively. In the following example (taken from section 7.1) the CLIPS rule is preceded by its linguistic equivalent.

If

a *subordinate-working-unit* has-superior an *unknown-superior*, and
 the *subordinate-working-unit* has-doctrinal-type a *subordinate-doctrinal-type*,
 and
 the *subordinate-doctrinal-type* has-kind a *doctrinal-kind*, and
 a *working-unit* has-doctrinal-type a *doctrinal-type*, and
 the *doctrinal-type* has-element an *element*, and
 the *element* has-doctrinal-type the *doctrinal-kind*,

Then

the *working-unit* is a candidate superior of the *subordinate-working-unit*.
 (Rule 1)

The actual CLIPS version of this rule includes a ranking of the candidate superior based on the euclidian distance separating the units.

```

(defrule determine-indirect-superior-of-unit-with-no-known-superior "and its
  ranking"
  (process determine-indirect-superior)
  (Logical-unit (serial ?LU1) (identifier ?identifier1) (status CURRENT))
  (Logical-unit-data (serial ?LU1) (type ?type1) (size ?size1) (location ?location1)
  (organic-superior UNKNOWN) (status CURRENT))
  (Logical-unit (serial ?LU2) (identifier ?identifier2) (status CURRENT))
  (Logical-unit-data (serial ?LU2) (type-for-resources ?type2) (size ?size2)
  (location ?location2) (status CURRENT))
  (DoctrinalType (doctrinal-type ?doctrinal-type2) (type ?type2) (size ?size2)
  (kind COMPOSITE))
  (Doctrine ?type1 has-kinks $? ?kind $?)
  (Doctrine ?doctrinal-type2 has-elements $? ?size1 $?)
  (Location (name ?location1) (xmin ?xmin1) (ymin ?ymin1) (xmax ?xmax1)
  (ymax ?ymax1))
  (Location (name ?location2) (xmin ?xmin2) (ymin ?ymin2) (xmax ?xmax2)
  (ymax ?ymax2))
  (test (spatially-intersectsp ?xmin1 ?ymin1 ?xmax1 ?ymax1 ?xmin2 ?ymin2
  ?xmax2 ?ymax2))
=>
(bind ?rank (spatial-distance ?xmin1 ?ymin1 ?xmax1 ?ymax1 ?xmin2 ?ymin2
?xmax2 ?ymax2))
(assert (IndividualQuery (subject ?identifier1) (relation element-of) (object
?identifier2) (rule determine-indirect-superior-of-unit-with-no-known-
superior) (rank ?rank) (ranking SPECIAL)))

```

(Rule 1 CLIPS version)

The Tactical Message Interpretation System

D. J. Parker

(DSTO-TR-0283)

DISTRIBUTION LIST

		Number of Copies
<i>Defence Science and Technology Organisation</i>		
Chief Defence Scientist and members of the)	1 shared copy
DSTO Central Office Executive)	for circulation
Counsellor, Defence Science, London		(Document Control sheet)
Counsellor, Defence Science, Washington		(Document Control sheet)
Senior Defence Scientific Adviser)	1 shared copy
Scientific Adviser - POLCOM)	
Director, Aeronautical & Maritime Research Laboratory		1
<i>Defence Intelligence Organisation</i>		
Assistant Secretary Scientific Analysis		1
<i>HQADF</i>		
Director Support Force Development (Land)		1
<i>Navy</i>		
Navy Scientific Adviser (NSA)		1
<i>Army</i>		
SO1 Int Dev, Land Headquarters, Land Int Branch		1
Scientific Adviser, Army (SA-A)		1
<i>Air Force</i>		
Air Force Scientific Adviser (AFSA)		1
<i>Electronics and Surveillance Research Laboratory</i>		
Chief, Information Technology Division		1
Research Leader, Command Control & Intelligence Systems		1
Research Leader, Military Computing Systems		1
Research Leader, Command Control and Communications		1
Head, Intelligence Systems Group		1
Head, Command Support Systems Group		1
Head, Information Management Group		1
Author [D. J. Parker]		1

DISTRIBUTION LIST

	Number of Copies
<i>Libraries and Information Services</i>	
Defence Central Library, Technical Reports Centre	1
Manager Document Exchange Centre (for retention)	1
National Technical Information Centre, United States	2
Defence Research Information Centre, United Kingdom	2
Director Scientific Information Services, Canada	1
Ministry of Defence, New Zealand	1
National Library of Australia	1
Defence Science and Technology Organisation Salisbury, Research Library	2
Library Defence Signals Directorate, Canberra	1
AGPS	1
British Library Document Supply Centre	1
Parliamentary Library of South Australia	1
The State Library of South Australia	1
<i>Spares</i>	
Defence Science and Technology Organisation Salisbury, Research Library	6

Department of Defence
DOCUMENT CONTROL DATA SHEET

1. Page Classification Unclassified
2. Privacy Marking/Caveat (of document) N/A

3a. AR Number AR-009-454	3b. Laboratory Number DSTO-TR-0283	3c. Type of Report Technical Report	4. Task Number ADL 93/134	
5. Document Date November 1995	6. Cost Code N/A	7. Security Classification	8. No of Pages	56
10. Title The Tactical Message Interpretation System		* <input type="checkbox"/> U <input type="checkbox"/> U <input type="checkbox"/> U Document Title Abstract S (Secret) C (Confi) R (Rest) U (Unclas)	9. No of Refs	13
11. Author(s) David J. Parker		12. Downgrading/Delimiting Instructions N/A		
13a. Corporate Author and Address Information Technology Division Electronics and Surveillance Research Laboratory PO Box 1500, Salisbury SA 5108		14. Officer/Position responsible for Security:CITD..... Downgrading:CITD.....		
13b. Task Sponsor DGFD(L)		Approval for Release:CITD.....		
15. Secondary Release Statement of this Document APPROVED FOR PUBLIC RELEASE				
16a. Deliberate Announcement No limitation				
16b. Casual Announcement (for citation in other documents) <input checked="" type="checkbox"/> No Limitation <input type="checkbox"/> Ref. by Author, Doc No. and date only				
17. DEFTEST Descriptors Command and control systems Information systems Military intelligence			18. DISCAT Subject Codes N/A	
19. Abstract This report describes an experimental computer support system for tactical military intelligence officers. The purpose of the system is to analyse structured military intelligence messages, and to derive conclusions similar to those that might be made by an expert intelligence officer. Details of the structure, functionality and input requirements of the system, and results using information taken from a military training exercise are presented.				